

**IN THE CLAIMS:**

Please write the claims to read as follows:

1 1. (Currently Amended) A method for a storage operating system implemented in a  
2 storage system to optimize an amount of readahead data retrieved from  
3 a data container of the storage system, ~~the method comprising:~~  
4 ~~maintaining a plurality of files;~~  
5 managing a separate set of readahead metadata for each individual read stream of  
6 a plurality of read streams, wherein the read stream is a set of one or more client read  
7 requests to retrieve data from a contiguous range of file offsets within a requested file;  
8 storing the separate set of readahead metadata in a corresponding individual  
9 readset data structure for the each individual read stream;  
10 ~~maintaining, for a selected file of the plurality of files, a plurality of in the~~  
11 individual readset data structures, each readset data structure holding a plurality  
12 of factors for a selected the each read stream, the plurality of factors determining  
13 ~~the system to adjust adaptively the amount of readahead data retrieved from the data~~  
14 container the one or more client read requests;  
15 receiving a client read request associated with for a particular read stream at the  
16 storage system and determining an amount of readahead data to retrieve from the data  
17 container for the particular read stream;  
18 receiving a next client read request associated with the particular read stream;  
19 locating a readset data structure for the particular read stream;  
20 determining whether the storage operating system is permitted to retrieve  
21 readahead data from the data container in response to the received next client read  
22 request; and  
23 if it is determined that the storage operating system is permitted to retrieve  
24 readahead data from the data container, performing the steps of,  
25 (i) modifying one or more of the plurality of factors stored within the readset data  
26 structure for the particular read stream;

27       (ii) adjusting, in response to modifying the one or more factors, an amount of  
28 readahead data to retrieve from the data container for the next client read request  
29 associated with the particular read stream, based on the plurality of factors stored within  
30 the readset data structure, wherein the amount of readahead data to retrieve for the next  
31 client read request is different from the amount of readahead data retrieved for the client  
32 read request;

33       (iii) retrieving the adjusted amount of readahead data from the data container; and

34       ~~(iii)~~iv) determining if the readset data structure meets a criteria for being updated,  
35 and if the readset data structure meets the criteria then updating the readset data structure.

1       2. (Original) The method of claim 1, wherein the data container is a file, directory, vdisk  
2 or lun.

1       3. (Original) The method of claim 1, wherein the storage operating system is determined  
2 to be permitted to retrieve readahead data from the data container when the client-  
3 requested data extends the read stream past a predetermined next readahead value.

1       4. (Original) The method of claim 3, wherein the predetermined next readahead value is  
2 stored in a readset data structure associated with the read stream.

1       5. (Previously Presented) The method of claim 3, wherein the predetermined next  
2 readahead value is updated based on a percentage of the adjusted amount of readahead  
3 data.

1       6. (Previously Presented) The method of claim 1, wherein a read-access style associated  
2 with the data container is one of the plurality of factors used to select the amount of  
3 readahead data.

1 7. (Previously Presented) The method of claim 6, wherein the adjusted amount of  
2 readahead data equals zero if the read-access style corresponds to a random read-access  
3 style.

1 8. (Previously Presented) The method of claim 1, wherein a number of client read  
2 requests processed in the read stream is one of the plurality of factors used to select the  
3 amount of readahead data.

1 9. (Original) The method of claim 8, wherein the number of client read requests  
2 processed in the read stream is stored as a count value in a readset data structure  
3 associated with the read stream.

1 10. (Previously Presented) The method of claim 1, wherein the amount of client-  
2 requested data is one of the plurality of factors used to select the amount of readahead  
3 data.

1 11. (Previously Presented) The method of claim 10, wherein the adjusted amount of  
2 readahead data is set equal to a predetermined upper limit for large amounts of client-  
3 requested data.

1 12. (Previously Presented) The method of claim 1, wherein the adjusted amount of  
2 readahead data is doubled if the number of client read requests processed in the read  
3 stream is greater than a first threshold value.

1 13. (Original) The method of claim 1, wherein the client-requested data is identified as  
2 read-once data when either (i) the number of client read requests processed in the read  
3 stream is greater than a second threshold value or (ii) a set of metadata associated with  
4 the read stream indicates that the client-requested data is read-once data.

1 14. (Previously Presented) The method of claim 1, wherein the adjusted amount of  
2 readahead data is stored in one or more buffers enqueued on a flush queue, the flush  
3 queue being configured to reuse buffers after a predetermined period of time.

1 15. (Original) The method of claim 14, wherein the predetermined period of time equals  
2 two seconds.

1 16. (Currently Amended) An apparatus configured to implement a storage operating  
2 system that optimizes an amount of readahead data retrieved from a data container of the  
3 apparatus, ~~the apparatus~~ comprising:

4 means for receiving a first data read command associated with a particular read  
5 stream, wherein the read stream is a set of one or more client read commands to retrieve  
6 data from a contiguous range of file offsets within a requested file;

7 ~~maintaining a plurality of files;~~

8 means for maintaining, for a selected file of ~~the~~ a plurality of files, a plurality of  
9 readset data structures, each readset data structure holding a plurality of factors for a  
10 ~~selected particular~~ read stream, the plurality of factors allowing the system to adjust  
11 adaptively the amount of data retrieved from the data container;

12 means for receiving a client read request associated with ~~for a the~~ particular read  
13 stream and means for determining an amount of readahead data to retrieve from the data  
14 container for the particular read stream;

15 means for receiving a next client read request associated with the particular read  
16 stream;

17 means for locating a readset data structure for the particular read stream;

18 ~~means for determining whether the storage operating system is permitted to~~  
19 ~~retrieve readahead data from the data container in response to the received client read~~  
20 ~~request;~~

21 means for adjusting an amount of readahead data to retrieve from the data  
22 container based on modifying one or more of the plurality of factors stored within the

23 | readset data structure, wherein the adjusted amount is a different amount than an amount  
24 | retrieved from the client read request;

25 |       means for retrieving the adjusted amount of readahead data from the data  
26 | container; and

27 |       means for determining if the ~~readset~~readahead data structure meets a criteria for  
28 | being updated, and if the ~~readset~~readahead data structure meets the criteria then updating  
29 | the ~~readset~~readahead data structure.

1 | 17. (Original) The apparatus of claim 16, wherein the data container is a file, directory,  
2 | vdisk or lun.

1 | 18. (Original) The apparatus of claim 16, wherein the storage operating system is  
2 | determined to be permitted to retrieve readahead data from the data container when the  
3 | client-requested data extends the read stream past a predetermined next readahead value.

1 | 19. (Previously Presented) The apparatus of claim 18, further comprising means for  
2 | updating the predetermined next readahead value based on a percentage of the adjusted  
3 | amount of readahead data.

1 | 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors  
2 | used to select the amount of readahead data includes at least one of:

- 3 |       (i) the amount of client-requested data,  
4 |       (ii) a number of client read requests processed in the read stream, and  
5 |       (iii) a read-access style associated with the data container.

1 | 21. (Previously Presented) The apparatus of claim 16, wherein the adjusted amount of  
2 | readahead data is doubled if the number of client read requests processed in the read  
3 | stream is greater than a first threshold value.

1 22. (Previously Presented) A storage system configured to optimize an amount of  
2 readahead data retrieved from a data container of the storage system, the storage system  
3 comprising:  
4 a network adapter for receiving a client read request for a particular read stream;  
5 a memory configured to store instructions for implementing a storage operating  
6 system that performs the steps of:  
7 locating a readset data structure for the particular read stream;  
8 adjusting, in response to the readset data structure for the particular read stream,  
9 the adjusted amount of readahead data from the data container based on a plurality of  
10 factors, the plurality of factors allowing the system to adjust adaptively the amount of  
11 data retrieved from the data container;  
12 retrieving the adaptively adjusted amount of readahead data from the data  
13 container.

1 23. (Original) The storage system of claim 22, wherein the data container is a file,  
2 directory, vdisk or lun.

1 24. (Original) The storage system of claim 22, wherein the storage operating system is  
2 determined to be permitted to retrieve readahead data from the data container when the  
3 client-requested data extends the read stream past a predetermined next readahead value.

1 25. (Previously Presented) The storage system of claim 24, wherein the predetermined  
2 next readahead value is updated based on a percentage of the adjusted amount of  
3 readahead data.

1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of  
2 factors used to select the amount of readahead data includes at least one of:  
3 (i) the amount of client-requested data,  
4 (ii) a number of client read requests processed in the read stream, and  
5 (iii) a read-access style associated with the data container.

1 27. (Previously Presented) The storage system of claim 22, wherein the adjusted amount  
2 of readahead data is doubled if the number of client read requests processed in the read  
3 stream is greater than a first threshold value.

1 28. (Currently Amended) A computer-readable media ~~containing~~comprising program  
2 instructions ~~executed for execution in~~ by a processor ~~for the practice of a method for a~~  
3 ~~storage operating system implemented in a storage system to optimize an amount of~~  
4 ~~readahead data retrieved from a data container of the storage system, the method~~  
5 comprising:

6 program instructions that receive a first data read command associated with a  
7 particular read stream, wherein the read stream is a set of one or more client read  
8 commands to retrieve data from a contiguous range of file offsets within a requested  
9 file~~maintaining a plurality of files~~;

10 program instructions that maintain~~maintaining~~, for a selected file of ~~the~~ a plurality  
11 of files, a plurality of readset data structures, each readset data structure holding a  
12 plurality of factors for a ~~selected particular~~ read stream based on a plurality of factors, the  
13 plurality of factors allowing the system to adjust adaptively the amount of data retrieved  
14 from the data container;

15 program instructions that receive~~receiving~~ a client read request associated with~~for~~  
16 a the particular read stream at the storage system and determining an amount of  
17 readahead data to retrieve from the data container for the particular read stream;

18 program instructions that receive a next client read request associated with the  
19 particular read stream;

20 program instructions that locate~~locating~~ a readset data structure for the particular  
21 read stream;

22 program instructions that determine~~determining~~ whether the storage operating  
23 system is permitted to retrieve readahead data from the data container in response to the  
24 received next client read request;

25           if it is determined that the storage operating system is permitted to retrieve  
26    readahead data from the data container, ~~performing executing program instructions~~  
27    ~~that the steps of:~~

28           (i) ~~adjusting~~ adjust an amount of readahead data to retrieve from the data  
29    container based on modifying one or more of the plurality of factors stored within the  
30    readset data structure; ~~and~~

31           (ii) ~~retrieving~~ retrieve the adjusted amount of readahead data from the data  
32    container; ~~and~~

33           (iii) ~~determining~~ determine if the readset data structure meets a criteria for being  
34    updated, and if the readset data structure meets the criteria then updating the readset data  
35    structure.

1    29. (Original) The computer-readable media of claim 28, wherein the data container is a  
2    file, directory, vdisk or lun.

1    30. (Previously Presented) The method of claim 1, wherein the retrieved readahead data  
2    is stored in one or more buffers, the buffers containing a flush queue, the flush queue  
3    being configured to reuse buffers after a predetermined period of time.

1    31. (Previously Presented) The method of claim 30, wherein the read stream  
2    corresponds to a read-once data transfer and data retrieved from the data container is  
3    stored in the flush queue.

1    32. (Previously Presented) The method of claim 30, wherein the retrieved readahead  
2    data is stored in the flush queue.

1    33. (Previously Presented) The method of claim 30, wherein one or more buffers  
2    accessed from the flush queue are re-enqueued on a normal queue.

1 34. (Currently Amended) A method for optimizing readahead data retrieved from  
2 a data container of a storage system, the method comprising:  
3 maintaining a plurality of files;  
4 maintaining, for a selected file of the plurality of files, a plurality of readset data  
5 structures, each readset data structure holding a plurality of factors for a ~~selected~~  
6 particular read\_stream, wherein the read stream is a set of one or more client read  
7 commands to retrieve data from a contiguous range of file offsets within a requested file;  
8 receiving a client read request for a particular read stream;  
9 receiving a next client read request associated with the particular read stream;  
10 locating a readset data structure for the particular read stream;  
11 adjusting an amount of readahead data based on ~~the~~an indicated amount of the  
12 next client-requested data and a corresponding readset data structure, wherein the amount  
13 of readahead data is adjusted in response to modifying on or more of the plurality of  
14 factors and wherein the adjusted amount is a different amount than an amount retrieved  
15 from the client read request; and  
16 retrieving the adjusted amount of readahead data from the data container.

1 35. (Previously Presented) The method of claim 34, wherein the adjusted amount of  
2 readahead data is set equal to a multiple of a predetermined amount, and wherein the  
3 multiple is associated with the amount of client-requested data.

1 36. (Previously Presented) The method of claim 34, wherein the adjusted amount of  
2 readahead data is set equal to a multiple of the amount of client-requested data.

1 37. (Previously Presented) The method of claim 36, further comprising the step of  
2 rounding the adjusted amount of readahead data to the size of a data block.

1 38. (Previously Presented) The method of claim 34, wherein the adjusted amount of  
2 readahead data is set equal to a predetermined upper limit.

1 39. (Currently Amended) A method for optimizing readahead data retrieved from a data  
2 container of a storage system, the method comprising:

3 maintaining a plurality of files;

4 maintaining, for a selected file of the plurality of files, a plurality of readset data  
5 structures, each readset data structure holding a plurality of factors for a selected  
6 readstream based on a plurality of factors, the plurality of factors allowing the system to  
7 adjust adaptively the amount of data retrieved from the data container;

8 receiving a client read request for a particular read stream at the storage system  
9 wherein the read stream is a set of one or more client read commands to retrieve data  
10 from a contiguous range of file offsets within a requested file;

11 locating a readset data structure for the particular read stream;

12 adjusting for the read stream an amount of readahead data based on a read-access  
13 style associated with the data container and modifying a corresponding readset data  
14 structure; and

15 retrieving the adjusted amount of readahead data from the data container.

1 40. (Previously Presented) The method of claim 39, wherein the adjusted amount of  
2 readahead data equals zero if the read-access style corresponds to a random read-access  
3 style.

1 41. (Currently Amended) A method for optimizing readahead data retrieval from  
2 a data container of a storage system associated with a number of storage devices, ~~the~~  
3 ~~method~~ comprising:

4 maintaining a plurality of files;

5 maintaining, for a selected file of the plurality of files, a plurality of readset data  
6 structures, each readset data structure holding a plurality of factors for a selected  
7 readstream based on a plurality of factors, the plurality of factors allowing the system to  
8 adjust adaptively the amount of data retrieved from the data container;

9 receiving a client read request for a particular read stream at the storage system  
10 wherein the read stream is a set of one or more client read commands to retrieve data  
11 from a contiguous range of file offsets within a requested file;

12 locating a readset data structure for the particular read stream;  
13 adjusting an amount of readahead data based on a number of storage devices and  
14 a corresponding readset data structure; and  
15 retrieving the adjusted amount of readahead data from the data container.

1 42. (Previously Presented) The method of claim 41, wherein adjusting an amount of  
2 readahead data further comprises:

3 determining whether a flag is associated with the read stream, the flag indicating  
4 that the storage system is associated with more than a predetermined number of storage  
5 devices; and

6 in response to determining whether the flag is associated, adjusting the amount of  
7 readahead data.

1 43. (Previously Presented) The method of claim 41, wherein the storage devices  
2 comprise one or more disks.

1 44. (Currently Amended) A method for optimizing readahead data retrieval in a storage  
2 system, ~~the method~~ comprising:

3 maintaining a plurality of files;  
4 maintaining, for a selected file of the plurality of files, a plurality of readset data  
5 structures, each readset data structure holding a plurality of factors for a selected  
6 readstream based on a plurality of factors, the plurality of factors allowing the system to  
7 adjust adaptively the amount of data retrieved from the data container;

8 receiving a client read request for a particular read stream at the storage system  
9 wherein the read stream is a set of one or more client read commands to retrieve data  
10 from a contiguous range of file offsets within a requested file;

11 locating a readset data structure for the particular read stream;

12 |        adjusting an amount of readahead data based on modifying one or more of a  
13 | plurality of factors stored within a corresponding readset data structure; and  
14 |        retrieving the adjusted amount of readahead data from a data container.

1    45. (Previously Presented) The method of claim 44, wherein the retrieved readahead  
2    data is stored in one or more buffers, the buffers containing a flush queue, the flush queue  
3    being configured to reuse buffers after a predetermined period of time.

1    46. (Previously Presented) The method of claim 45, wherein the read stream  
2    corresponds to a read-once data transfer and data retrieved from the data container is  
3    stored in the flush queue.

1    47. (Previously Presented) The method of claim 45, wherein the retrieved readahead  
2    data is stored in the flush queue.

1    48. (Previously Presented) The method of claim 45, wherein one or more buffers  
2    accessed from the flush queue are re-enqueued on a normal queue.

1    49. (Currently Amended) A computer system for optimizing readahead data retrieval in  
2    a computer data storage system, ~~the system~~ comprising:

3        maintaining a plurality of files;

4        maintaining, for a selected file of the plurality of files, a plurality of readset data  
5    structures, each readset data structure holding a plurality of factors for a selected  
6    readstream based on a plurality of factors, the plurality of factors allowing the system to  
7    adjust adaptively the amount of data retrieved from the data container and wherein the  
8    read stream is a set of one or more client read commands to retrieve data from a  
9    contiguous range of file offsets within a requested file;

10       means for receiving a client read request for a particular read stream at the storage  
11    system;

12       means for locating a readset data structure for the particular read stream;

13 | means for adjusting an amount of readahead data based on modifying one or more  
14 | of a plurality of factors stored within a corresponding readset data structure; and  
15 | means for retrieving the adjusted amount of readahead data from a data container.

1 50. (Previously Presented) The system of claim 49, wherein the retrieved readahead data  
2 is stored in one or more buffers, the buffers containing a flush queue, the flush queue  
3 being configured to reuse buffers after a predetermined period of time.

1 51. (Previously Presented) The system of claim 50, wherein the read stream corresponds  
2 to a read-once data transfer and data retrieved from the data container is stored in the  
3 flush queue.

1 52. (Previously Presented) The system of claim 50, wherein the retrieved readahead data  
2 is stored in the flush queue.

1 53. (Previously Presented) The system of claim 50, wherein one or more buffers  
2 accessed from the flush queue are re-enqueued on a normal queue.

1 54. (Currently Amended) A method, comprising:  
2 maintaining a plurality of files;  
3 maintaining, for a selected file of the plurality of files, a plurality of readset data  
4 structures, each readset data structure holding a plurality of factors for a selected  
5 readstream based on a plurality of factors, the plurality of factors allowing the system to  
6 adjust adaptively the amount of data retrieved from the data container and wherein the  
7 read stream is a set of one or more client read commands to retrieve data from a  
8 contiguous range of file offsets within a requested file;  
9 receiving a plurality of client read requests for a particular read stream at a storage  
10 system;  
11 locating a readset data structure for the particular read stream;

12           adjusting an amount of readahead data to retrieve from one or more data  
13 containers based on a plurality of factors stored within the readset data structure;  
14           retrieving the adjusted amount of readahead data from the data container;  
15           processing one or more of the plurality of client read requests; and  
16           adjusting, as client requests are processed, the plurality of factors stored within  
17 the readset data structure associated with each read stream, whereby adjusting the  
18 plurality of factors further adjusts to optimize the amount of readahead data that is cached  
19 for each client read request directed to the particular read stream.

1   55. (Previously Presented) The method of claim 54, further comprising:  
2           determining whether the storage operating system is permitted to retrieve  
3 readahead data from the one or more data containers in response to each received client  
4 read request.

1   56. (Previously Presented) The method of claim 54, wherein the one or more data  
2 containers are at least one of a file, a directory, a vdisk or a lun.

1   57. (Previously Presented) The method of claim 55, wherein the storage operating  
2 system is determined to be permitted to retrieve readahead data from the one or more data  
3 containers when the client-requested data extends the read stream past a predetermined  
4 next readahead value.

1   58. (Previously Presented) The method of claim 57, wherein the predetermined next  
2 readahead value is stored in a readset data structure associated with the read stream.

1   59. (Previously Presented) The method of claim 57, wherein the predetermined next  
2 readahead value is updated based on a percentage of the adjusted amount of readahead  
3 data.

1 60. (Previously Presented) The method of claim 54, wherein a read-access style  
2 associated with the one or more data containers is one of the plurality of factors used to  
3 select the amount of readahead data.

1 61. (Previously Presented) The method of claim 60, wherein the adjusted amount of  
2 readahead data equals zero if the read-access style corresponds to a random read-access  
3 style.

1 62. (Previously Presented) The method of claim 54, wherein a number of client read  
2 requests processed in the read stream is one of the plurality of factors used to select the  
3 amount of readahead data.

1 63. (Previously Presented) The method of claim 62, wherein the number of client read  
2 requests processed in the read stream is stored as a count value in a readset data structure  
3 associated with the read stream.

1 64. (Previously Presented) The method of claim 54, wherein the amount of client-  
2 requested data is one of the plurality of factors used to select the amount of readahead  
3 data.

1 65. (Previously Presented) The method of claim 64, wherein the adjusted amount of  
2 readahead data is set equal to a predetermined upper limit for large amounts of client-  
3 requested data.

1 66. (Previously Presented) The method of claim 54, wherein the adjusted amount of  
2 readahead data is doubled if the number of client read requests processed in the read  
3 stream is greater than a first threshold value.

1 67. (Previously Presented) The method of claim 55, wherein the client-requested data is  
2 identified as read-once data when either (i) the number of client read requests processed

3 in the read stream is greater than a second threshold value or (ii) a set of metadata  
4 associated with the read stream indicates that the client-requested data is read-once data.

1 68. (Previously Presented) The method of claim 54, wherein the adjusted amount of  
2 readahead data is stored in one or more buffers enqueued on a flush queue, the flush  
3 queue being configured to reuse buffers after a predetermined period of time.

1 69. (Currently Amended) A method for optimizing readahead data retrieval for a storage  
2 system, ~~the method~~ comprising:

3 maintaining a plurality of files;

4 maintaining, for a selected file of the plurality of files, a plurality of readset data  
5 structures, each readset data structure holding a plurality of factors for a selected read  
6 stream, wherein the read stream is a set of one or more client read commands to retrieve  
7 data from a contiguous range of file offsets within a requested file;

8 receiving a client read request for a particular read stream at the storage system;

9 locating a readset data structure for the particular read stream;

10 adjusting an amount of readahead data in response to a corresponding readset data  
11 structure based on modifying one or more of a plurality of factors, the plurality of factors  
12 allowing the system to adjust adaptively the amount of data requested to be retrieved  
13 from the data container by one or more client read requests associated with the particular  
14 read stream; and

15 retrieving the adaptively adjusted amount of readahead data from a data container.

1 70. (Previously Presented) The method of claim 69, further comprising:

2 allocating more readsets for the file in response to processing one or more client  
3 “write” requests to the file.

1 71. (Currently Amended) A storage system, comprising:

2 an operating system to maintain a plurality of files, and to maintain, for a selected  
3 file of the plurality of files, a plurality of readset data structures, each readset data

4 | structure holding a plurality of factors for a selected read\_stream, wherein the read stream  
5 | is a set of one or more client read commands to retrieve data from a contiguous range of  
6 | file offsets within a requested file;

7 | a network adapter to receive a client read request at the storage system for client-  
8 | requested data stored in a file;

9 | a plurality of readset data structures associated with the client-requested data;

10 | an operating system to establish a read stream corresponding to each readset data  
11 | structure of the plurality of readset data structures;

12 | a process to adjust an amount of readahead data in response to a corresponding  
13 | readset data structure of the plurality of readset data structures which is based on a  
14 | plurality of factors, the plurality of factors allowing the system to adjust adaptively the  
15 | amount of data requested to be retrieved, by one or more client read requests directed to  
16 | the particular read stream, from the data container;

17 | a storage adapter to retrieve the adjusted amount of readahead data from the data  
18 | container; and

19 | the operating system to determine if the readset data structure meets a criteria for  
20 | being updated, and if the readset data structure meets the criteria, then updating the  
21 | readset data structure.

1 | 72. (Previously Presented) The storage system of claim 71, further comprising:

2 | the operating system to allocate more readsets for the file in response to  
3 | processing one or more client “write” requests to the file.

Please add new claims 73, *et seq.* as follows:

- 1    73. (New) A method for operating a computer data storage system, comprising:  
2            receiving a first data read command associated with a particular read stream,  
3    wherein the read stream is a set of one or more client read commands to retrieve data  
4    from a contiguous range of file offsets within a requested file;  
5            determining one or more input parameters of the first data read command;  
6            establishing, in response to the input parameters, a readahead hint, wherein the  
7    readahead hint determines a number of data blocks to readahead;  
8            receiving a next data read command associated with the particular read  
9    stream;  
10           determining one or more input parameters of the next data read command;  
11           modifying, in response to the next input parameters, the readahead hint to  
12    obtain a modified readahead hint; and  
13           adjusting, in response to the modified readahead hint, the number of data  
14    blocks to readahead for the next read command associated with the particular read  
15    stream, wherein the adjusted number of data blocks is stored in memory.
- 1    74. (New) The method of claim 73 further comprising:  
2            modifying the readahead hint by writing the input parameter into an associated  
3    readahead metadata of the read stream.
- 1    75. (New) A method for operating a computer data storage system, comprising:  
2            receiving a first data read command associated with a particular read stream,  
3    wherein the read stream is a set of one or more client read commands to retrieve data  
4    from a contiguous range of file offsets within a requested file;  
5            determining one or more parameters associated with the first data read command;  
6            storing the parameters in a readahead metadata associated with the particular read  
7    stream, wherein the metadata determines a number of data blocks to readahead;  
8            receiving a next data read command associated with the particular read stream;  
9            determining one or more parameters of the next data read command;

10           modifying the metadata in response to the parameters of the next data read  
11   command; and  
12           adjusting, in response to the modified metadata, the number of data blocks to  
13   readahead for the next read command associated with the particular read stream, wherein  
14   the adjusted number of data blocks is stored in memory.